Articles

Timing Is Everything
Time-Oriented Clinical Information Systems

YUVAL SHAHAR, MD, PhD, *Stanford, California, and* CARLO COMBI, PhD, *Udine, Italy*

Time is important in clinical information systems. Representing, maintaining, querying, and reasoning about time-oriented clinical data is a major theoretical and practical research area in medical informatics. In this nonexhaustive overview, we present a brief synopsis of research efforts in designing and developing time-oriented information systems in medicine. These efforts can be viewed from either an application point of view, distinguishing between different clinical tasks (such as diagnosis versus therapy) and clinical areas (such as infectious diseases versus oncology), or a methodological point of view, distinguishing between different theoretical approaches. We also explore the two primary methodological and theoretical paths research has taken in the past decade: temporal reasoning and temporal data maintenance. Both of these research areas include efforts to model time, temporal entities, and temporal queries. Collaboration between the two areas is possible, through tasks such as the abstraction of raw time-oriented clinical data into higher-level meaningful clinical concepts and the management of different levels of temporal granularity. Such collaboration could provide a common ground and useful areas for future research and development. We conclude with our view of future research directions.

(Shahar Y, Combi C. Timing is everything—time-oriented clinical information systems. West J Med 1998; 168:105–113)

It is almost inconceivable to represent and reason about clinical data without a temporal dimension. Clinical interventions must occur at one or more time points (a tonsillectomy performed *on January 23, 1992, at 9:00 AM*); contain certain facts such as laboratory test results or a diagnosis; show duration of time periods (moderate anemia known at least *from March 15, 1997, to August 19, 1997*); and present temporal relationships (high fever occurring *after* mumps immunization). In medical information systems, the element of time is instrumental in performing several functions: representing information within a computer-based electronic medical-record system; querying medical records; and reasoning about time-oriented clinical data as part of various decision-support applications, such as diagnosis, therapy, and the general browsing of electronic patient records. The proper performance of these functions is equally important for health care professionals who need certain information about one or more patient records for automated decision-support systems. For instance, while treating a patient with an experimental chemotherapy protocol, either a health care professional or an intelligent therapy–support system may need to know the number of episodes of bone marrow toxicity, as well as the severity, length, and last occurrence of each episode.

It useful to distinguish between two research directions, which can be found in both the computer-science and medical-informatics literature (Figure 1). Each direction is distinct with respect to its focus and the research communities pursuing it. *Temporal reasoning* involves various inference tasks using time-oriented clinical data, such as therapy planning and execution; it has traditionally been linked with the artificial-intelligence community. *Temporal data maintenance*, which deals with the storage and retrieval of clinical data that have heterogeneous temporal dimensions, is typically associated with the (temporal) database community. A theme common to both directions is the necessity for *temporal data modeling*, without which clinical data can be neither maintained nor reasoned with.

A wide variety of applications use the temporal aspects of clinical data. Examples include the management of time-oriented data stored in the medical records of ambulatory or hospitalized patients,[1–6] the prediction of future values of clinical data given past trends,[7–9] the abstraction of time-oriented clinical data,[10–12] and the knowledge-based support of health

From the Section on Medical Informatics, School of Medicine, Stanford University, Stanford, California; and the Dipartimento di Matematica e Informatica, Universitá degli Studi di Udine, Udine, Italy.

Reprint requests to Yuval Shahar, MD, PhD, Section on Medical Informatics, Medical School Office Building x215, 251 Campus Drive, Stanford University, Stanford, CA 94305-5479. E-mail shahar@smi.stanford.edu.

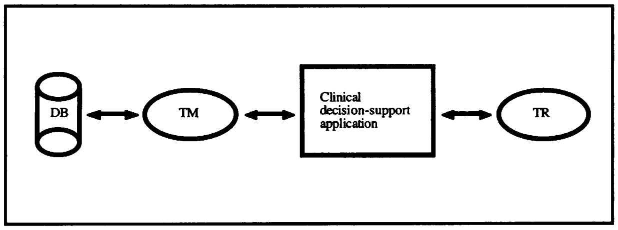

Figure 1.—This diagram illustrates the typical relationship of time-oriented computational modules in a clinical decision-support system. DB=patient electronic database; TM=temporal data-maintenance module; TR=temporal-reasoning module

care professionals' decisions based on time-oriented clinical data, such as diagnosis support, patient monitoring, or therapy planning.[13–17]

Studies of time-oriented applications have been performed in multiple clinical areas: cardiology,[3,4,18,19–21] oncology,[6,10,15] psychiatry,[22] internal medicine,[11,13,23] intensive care,[7,15,24,25] cardiac surgery,[26] orthopedics,[14] urology,[27] infectious diseases,[6] anesthesiology,[8,24] pediatrics,[21] and endocrinology.[28] Various clinical tasks are supported in an automated or semiautomated manner by the software systems proposed in the literature: diagnosis,[14] therapy administration and monitoring,[10,24,25] protocol- and guideline-based therapy,[6,11,17] and patient management.[1,4,29]

This article describes the main features characterizing the medical-informatics research (most of which includes implemented software systems at various stages of deployment) that deals with time-oriented clinical systems. We describe the problems related to the modeling of time-oriented clinical concepts; we then provide a brief overview of the research in temporal reasoning in medicine and discuss temporal maintenance in clinical information systems. These tasks form a potential bridge between the work on temporal data-maintenance systems and the work on temporal reasoning systems—in particular, time-oriented clinical decision support systems.

Clinical databases and the modeling of temporal concepts

A common focus of research regarding time-oriented clinical data is the definition or adoption of a set of basic concepts that describe the time-oriented clinical world in a sound and unambiguous way. Within medical informatics, this effort has progressed from an ad-hoc definition of concepts supporting a particular application to the adoption and proposal of more generic definitions, supporting different clinical applications.[4,6,11,14–16,26,30,31] For example, early work in interpreting real-time quantitative data in the intensive-care domain emphasized the problems of a module that suggests the optimal ventilator therapy at a given time.[15] A more recent framework,[16] however, uses a generic temporal ontology (a set of terms and concepts and the relations between them) and a general, comprehensive model of diagnostic reasoning.

Most clinical and research databases are currently known as *relational databases*. In relational databases, patient records are organized as a set of tables, where each row of a table (also called a *relation* or *tuple*) represents a flat record (such as a patient's identification, a visit date, a laboratory test name, or a laboratory test result). Queries are directed to the relational database using a well-defined algebraic model that uses a small amount of understood operations, such as, "*select* from the VISITS file all tuples in which the value of the VISIT-DATE field is between 10/1/96 and 9/31/97."

An alternative database in use today is known as the *object database*, in which important concepts—such as patients, visits, and laboratory data—are modeled as a hierarchy of independent objects with complex links (interrelations). Such a system allows flexibility in modeling and querying, but the semantics of the formal query language used is often less clear.

Several related concepts involving time appear in the medical-informatics literature. We distinguish between two related issues: modeling the concept of time and modeling the entities that have a temporal dimension.

Modeling time

There are three basic choices have to be made when modeling time for the management of or reasoning about time-oriented clinical data.

Time instants versus time intervals. The first choice involves establishing the difference between *time instants* (relating to instantaneous events such as a myocardial infarction) and *time intervals* (relating to situations that are interval-based, or lasting for a span of time). Both concepts have been used in the medical informatics literature to represent time.[4,11,16,26] Care must be taken in associating these concepts with clinical entities such as symptoms, therapies, and pathologies: a myocardial infarction, for example, could be considered either an instantaneous event, within the overall clinical history of the patient, or an interval-based concept, if observed, perhaps, during a stay in the ICU.

Another difference exists between the basic time primitives, usually instants (time points), and the basic time entities that can be associated with clinical concepts.[4,6,11] In defining basic time entities, time points are often adopted. Intervals are then represented by their upper and lower temporal bounds (start and end time points). Most systems employed in medical informatics applications use time points, rather than time intervals, as the basic time primitives—an approach that originated from research in artificial intelligence.[32] Several variations exist. One approach inspired by the artificial-intelligence area is to use a set of *time stamps* as the basic time primitives, from which both time points and time intervals can be constructed. Assertions, such as values of clinical parameters or diagnoses, can only hold over time intervals, which are defined as ordered pairs of time stamps (time points are thus zero-length intervals).[33]

The first approach to associating time with clinical entities therefore explicitly includes both instant-related and interval-related entities.[26] The second approach associates only clinical entities with a certain type of time concept—usually an interval—dealing, in a homogeneous way, with intervals that degenerate to a single instant.[4,6,11]

Linear, branching, and circular times. Different properties can be associated with a time axis comprising instants: in both general and clinically-oriented databases, time is usually linear—that is, the set of time points is completely ordered.[4,6,16] For diagnosis, projection, or forecasting (such as predicting a clinical evolution over time), however, a branching time might be necessary. This method has been found to be useful, for instance, when describing potential outcomes in a pharmacoeconomics clinical trial, and it has been implemented using an object-oriented temporal model.[34] Circular time is used to describe recurrent events, such as "the administration of regular insulin every morning"; it is useful, for instance, when specifying procedural therapy guidelines and their underlying intentions.[35]

Absolute and relative times. The position on the time axis of an interval or of an instant can be given as an *absolute* position, such as in calendaric time[1,6,11,23,36] ("tachycardia on November 3, 1996"). This is a common approach adopted by temporal clinical databases. However, it also is common in medicine to use *relative* time references, such as "angina after a long walk" or "several episodes of headache during puberty." Incorporation of purely relative time-oriented, interval-based information (especially with disjunctions, such as "the patient had vomited before or during the episode of diarrhea") within a standard temporal database is still a difficult task, as we will discuss in the section on temporal granularity and uncertainty.

Modeling temporal entities

Two questions that have been investigated in some depth in the medical informatics literature are, What are the basic medical concepts that have temporal dimension? and, How should time-oriented clinical data be modeled? In general, we find two approaches in modeling temporal entities in medical applications: addition of a temporal dimension to existing objects, and creation of task-specific, time-oriented entities.

The first approach, originating from research on databases, adds one or more temporal dimensions to the basic existing medical-record entities, such as lab tests and clinical characterizations.[6,37] Within this approach, both instant- and interval-based information can be handled homogeneously, using the concept of a *temporal assertion*.[4] A temporal dimension can also be added to a relational database tuple, in a temporal extension of the database relational model, to aid clinical databases used for decision-support application.[6] In such approaches, complex temporal features of clinical data can be questioned by suitable query languages.[6,37] Indeed, one of the first applications of databases to clinical domains, explicitly addressing the time representation problem, was the Time-Oriented Database model.[1] This model has been adopted by the American Rheumatism Association Medical Information System to manage data related to the long-term clinical courses of patients suffering from arthritis or, more generally, from rheumatic pathologies.[38] The Time-Oriented Database uses a "cubic" vision of clinical data, meaning that values of patient visit data are indexed by three identifiers: patient identification number, time (visit date), and clinical-parameter type. Specialized time-oriented queries enable researchers to extract, for particular patients, data values that follow simple temporal patterns (such as an increase of a given rate). Assigning a temporal dimension to the tuple level is common to many applications of clinical databases.[3]

The second approach originates mostly from the area of artificial intelligence in medicine. It focuses on modeling different temporal features of complex, task-specific entities. The entities are defined by the needs of relevant temporal-reasoning tasks. (see Section 3). Based on the temporal entities that are stored at the database level, several types of compound (abstract) entities are introduced.

For example, in the HyperLipid system—used in the management of hypercholesterolemia[13]—patient visits were modeled as instant-based objects called *events*, while drug administration was modeled as objects called *therapy*, the attributes of which included a time interval. *Phases* of therapy (inspired by the systematically modeled clinical algorithm) were then introduced to model groups of heterogeneous data related to both visits and therapies. The events, therapies, and phases were then connected through a network.

A more general and influential model than that of the HyperLipid system is the *temporal network* model.[31] In the temporal-network model, a T-node (an object that contains temporal information within a temporal network) models task-specific temporal data, such as a chemotherapy cycle, at different levels of abstraction. Each T-node has a time interval, during which the information represented by the T-node's data is true for a given patient. The M-HTP system, used to monitor heart-transplant patients,[26] structures a patient's clinical facts in a temporal network. Through this network, a physician can obtain various temporal views of the patient's clinical history. Each node of the temporal network represents either an *event* (a patient visit) or a *significant episode* in the patient's clinical record. An event is time-point based; its temporal location can be specified by an absolute date or by the temporal distance to the transplantation event. A significant episode occurs during an interval, in which a predefined property (evaluated by reasoning about several events) is true. In another system used to represent skeletal dysplasias, the concepts of *findings*, *features*, and *events* were introduced to distinguish various types of instantaneous and interval-based information, either patient-specific or general.[14]

Temporal Reasoning

Temporal reasoning has been used in medical domains as part of a wide variety of generic tasks, such as projecting, forecasting, planning, interpreting, and diagnosing—specifically in temporal abstraction and monitoring. These tasks are often interdependent. *Projecting* involves determining the likely consequences of a set of conditions or actions, usually given as a set of cause–effect relationships. Projecting is particularly relevant to the planning task; it is used to determine, for example, the subsequent state of a patient when a drug with known side effects is administered. *Forecasting* includes predicting future values for various parameters given a vector of time-stamped past and present measured values. *Planning* consists of producing a sequence of actions for a health care professional, given the initial state of the patient, to achieve a goal state(s). The actions taken are usually operators with predefined certain or possible effects on the environment, and, to be possible or effective, they might require a set of enabling *preconditions*. Achieving the goal state, as well as some of the preconditions, might depend on correct projection of the actions up to a point, in order to determine whether the preconditions will hold when required. *Interpreting* involves the abstraction of a set of time-oriented patient data, either to an intermediate level of meaningful temporal patterns—as is common in the *temporal abstraction* task or the *monitoring* task—or to the level of a definite diagnosis or set of diagnoses that explain findings and symptoms, as is common in the *diagnosing* task. Interpreting, unlike forecasting and projecting, involves reasoning about past and present data and not about the future.

From a methodological point of view, a criterion that can be used when classifying temporal-reasoning research that has been applied to clinical data is whether it uses a deterministic or a probabilistic approach. Deterministic approaches are based on either well-known formalisms from the artificial-intelligence field or ad-hoc rules or ontologies.[16] Probabilistic approaches typically are associated with the tasks of interpreting or forecasting time-stamped clinical data whose values are affected by different sources of uncertainty.[9] A promising recent approach is the Dynamic Network Models methodology, a synthesis of belief-network (Bayesian-network) and classic time-series models. This methodology has been applied with encouraging results to tasks such as predicting the outcomes of critically ill intensive-care patients[7] and episodes of apnea in sleep-apnea patients.[8]

A task that needs to be commonly performed is the *temporal abstraction task*: a task that involves the abstraction of high-level concepts (such as a pattern of bone-marrow toxicity specific to a particular chemotherapy-related context) from time-oriented clinical data (such as a time-stamped series of chemotherapy-administration events and various hematological laboratory tests).[10,11,14,26,31] Several solutions have been proposed to the recurring problem of performance of the task. For example, in the M-HTP system for monitoring heart transplant patients[26] the "white blood-cell [WBC] count" measured during a visit occurs as an instantaneous event in the knowledge base, indexed by the visit date; "WBC-count decrease" is an episode, spanning several days, detected by the values of the WBC count. Similarly, in a temporal-network model,[31] T-nodes are able to describe, at different levels of abstraction, data related to a patient undergoing different chemotherapy treatments. In both instances, a temporal knowledge–based reasoner that uses what are known as *IF-THEN rules* is applied to the system's temporal model of the patient.[26,31] These systems are able to deal with complex temporal conditions. A typical rule, for example, is:[12]

IF DURING last 10 days ARE PRESENT
low CMV antigenemia OF TIME SPAN at least 7 days
AND
leukopenia OF TIME SPAN at least 5 days
AND
DURING last 15 days IS NOT PRESENT
CMV infection OF TIME SPAN at least 1 day
THEN CMV infection is highly suspected

An example of a diagnostic-support system is the Skeletal Dysplasia Diagnosis (SDD) expert system and its temporal reasoning framework.[14] The SDD system is designed to aid in diagnosing skeletal dysplasias and syndromes. The temporal reasoner module in the SDD system has a layered architecture and is able to accommodate findings and features, which, in turn, provides the user with higher-level representations of findings and dysplasia expectations for the patient.

A common, highly useful task that involves some amount of abstraction and a certain amount of forecasting, is the *validation* of time-oriented clinical data. This task typically requires both deterministic and probabilistic techniques, and possibly includes the suggestion of specific repairs to the data. Recent work[25] proposes several useful methods for validating and repairing high-frequency time-oriented clinical data; these methods have been applied to the neonatal intensive care area—specifically to artificial-ventilation.

Maintenance of Time-Oriented Clinical Data — Temporal Databases

Regarding the task of managing time-oriented clinical data, we observe that the literature has progressed from the early systems, which were mostly application dependent, to recent, more general approaches, which, even when applied to the solution of real problems in management of time-oriented clinical data, have a more generalized value and inherent soundness.[1,3,4,6,29,37,39–41] At first, systems that were designed to manage temporal clinical data were based on the flat relational-database model.[1,39] These systems were based on time-stamping the database tuples: the date of the visit was added to the specific attribute values. Later work[31] has proposed the use of a specific temporal-query language for clinical data that are structured by a temporal-network model. Although that

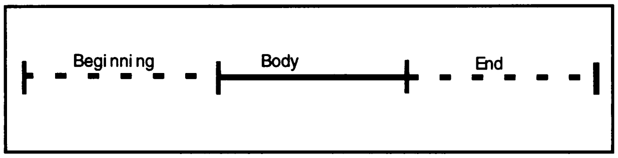

Figure 2.—This continuum shows a variable time interval. Variable intervals are composed of a certain body, during which the relevant property is sure to hold, and of uncertain start and end points, represented as intervals of uncertainty.

language was patient oriented and not based on a generic data model, it was one of the first proposals for an extension of query languages that would enable the system to retrieve complex temporal properties of stored data. Because most query languages and data models used for clinical data management were application-dependent, developers had to provide ad-hoc facilities for querying and manipulating specific temporal aspects of data.[31]

Recent work on temporal clinical databases presents a more general approach. Extensions of common data models, particularly of relational models, are also based on the general database field literature, which has given special attention to temporal databases in the past few years. One issue that has been explored in depth concerned what *kinds* of temporal dimensions must be supported by the temporal database. Three different temporal dimensions have been distinguished.[42] The first is the *transaction time*, that is, the time at which data are stored in the database: for example, the time in which the assertion "WBC count is 7600" was entered into the patient's medical record. The second dimension is the *valid time*, the time at which the data are true for the modeled real world entity: for example, the time in which the WBC-count was, in fact, 7600. The third dimension is the *user-defined time*, the meaning of which is related to the application and is thus defined by the user: for example, the time in which the WBC count was determined in the laboratory.

Using this temporal-dimension taxonomy, four kinds of databases can be defined: *snapshot databases*, based on flat, timeless data models; *rollback databases*, which explicitly represent only the transaction time (such as a series of updates to the patient's current address stamped by the time in which each modification was recorded); *historical databases*, which explicitly represent only the valid time (for instance, the most current knowledge about the WBC value on 1/12/97, allowing future updates referring to data on 1/12/97 but keeping no record of the updates themselves); and what are now called *bitemporal databases*, which explicitly represent both transaction time and valid time, making them both historical and rollback. Thus, in a bitemporal database we can explicitly represent that, on January 17, 1997 (transaction time), the physician entered in the patient's record the fact that on January 12, 1997 (valid time), the patient had an allergic reaction to a sulpha-type drug. There are many advantages to the use of bitemporal databases in medical information systems, including the ability to answer both research-related and legal questions such as, "When another physician prescribed sulpha on January 14, 1997, did she know at that time that the patient had an allergic reaction to sulpha on a previous date?"

In medical informatics, attention has been paid mostly to historical databases, which emphasize valid time and extend relational or object-oriented models.[6,37,40,48] For example, researchers have defined four different types of relational tuples—event, start, body, and stop—to specify, respectively, instantaneous facts and three aspects of uncertainty about interval-based facts: uncertainty about the start time of the fact, the end time of the fact, and a certain period of time in which the fact existed.[6] The resulting time interval is often called a *variable time interval*,[40] and the uncertainty regarding the start or the end of that interval is sometimes referred to as the *interval of uncertainty*[6] (Figure 2). The framework can be completed by an extension of the relational-database algebraic formula to manage temporal information and temporal relational operations.[6]

A goal of the interval of uncertainty representation is to facilitate the handling of temporal uncertainty and, in particular, the management of data represented at variable levels of temporal granularity, a task that we discuss further below.

Other researchers[37] have extended an object-oriented data model and its related query language to deal with temporal clinical data, taking into account different and mixed temporal granularities. Recent work has adapted an existing object database model to the management of time-oriented data and applied it to the modeling of pharmacoeconomic clinical trials.[34] The broad set of types supported by the adopted object data model enables a modeling of branching timelines; these correspond, for instance, to the evaluation of different pharmacological treatments.

Temporal Abstraction and Management of Temporal Granularity

Two commonly recurring and closely related tasks in both the temporal-reasoning and the temporal data–maintenance research areas are the temporal-abstraction task and the management of variable temporal granularity, both of which are discussed above. Since these tasks are relevant to both research communities, they might be viewed as a potential bridge between them (other bridges include fundamental issues, such as the time model, also mentioned above). The two tasks have been investigated in the fields of both medical informatics and general computer-science.

Temporal abstraction provides a more powerful, concise, and integrated description of a collection of time-stamped raw data (Figure 3). The term "temporal abstraction," however, is somewhat misleading, because it is the time-oriented data, and not the time itself, that are being abstracted. In the medical-informatics field, temporal abstraction plays a central role in supplying health care professionals with data at a level suitable for decision-making support. Temporal abstraction in general, and in medicine in particular, has been heavily investigated in recent years.[10–12,21,26,33]

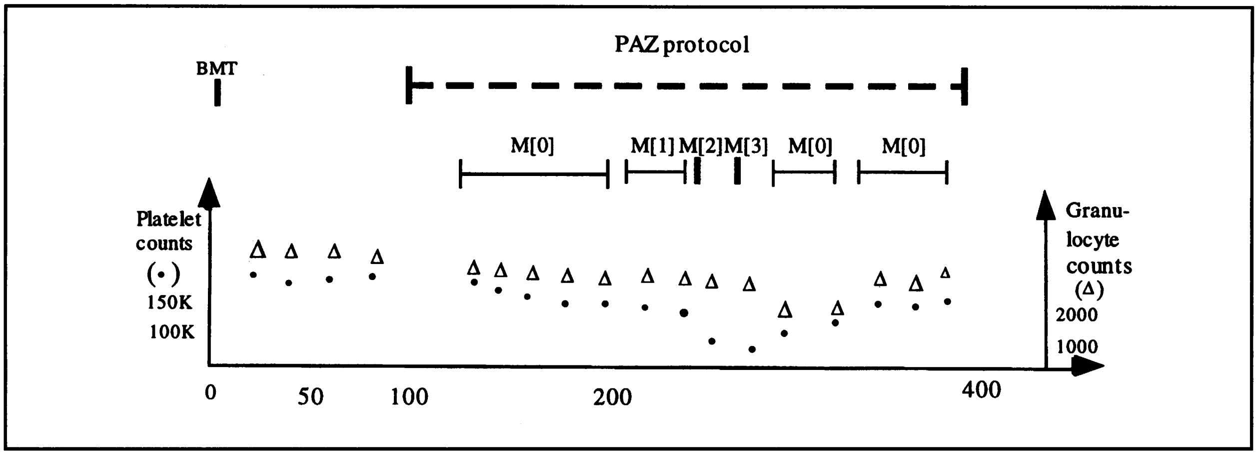

Figure 3.—An example of temporal abstraction in the domain of bone-marrow transplantation. The figure presents context-sensitive abstractions of platelet and granulocyte values into bone-marrow toxicity intervals during administration of the prednisone/azathioprine (PAZ) protocol for treating patients who have chronic graft-versus-host disease following a bone-marrow transplantation (BMT) event. • = platelet counts; △ = granulocyte counts; – – – – = event; —— = abstraction interval; M[*n*] = myelotoxicity (bone-marrow–toxicity) grade *n*, as defined in the context of PAZ therapy.

Management of variable temporal granularity deals with an abstraction of the time primitives themselves; it concerns the level of abstraction or the time unit (such as a day or a month) at which the time element (instant, interval, and so on) associated with the relevant data is represented. Using this definition, we observe that the tasks of temporal abstraction and the management of variable temporal granularities are interconnected. When reasoning about various temporal-granularity levels, emphasis is placed on abstracting the representation of the *time component* of a time-oriented assertion. When performing a temporal-abstraction task, however, emphasis is placed on the abstraction of the time-oriented *entity* itself.

There are three main types of temporal granularity:[4]

Abstraction granularity. This granularity-management aspect is not directly related to the time axis. Abstraction granularity refers to the ability to express complex and composite temporal concepts: "An anemia level that is increasing during a period of three months."

Absolute-time granularity. This is the ability to express the temporal dimension of the data by mixing and using different absolute time references. Absolute-time granularity refers to the uncertainty in specifying a temporal dimension or to the use of different time-units: "The vomiting episode began somewhere within the time interval starting on January 21, at 15:23, and ending on January 21, at 16:34."

Calendar-date granularity. This is the capacity for expressing the temporal dimension through the use of multiple time units, such as years, months, and days: "The diarrhea started during February 1997."

Abstraction granularity

Medical decision-support systems often do not associate the granularity of time with calendar time. Rather, the temporal granularity level is affected by the abstraction needed by the relevant clinical problem.[14,26,31,43]

Many representations of temporal data at high abstraction levels in medical expert systems were inspired by Allen's interval-based logic.[32] The temporal-network model aimed to extend the Time Oriented Database model by defining suitable persistent objects.[31] As we point out above, a temporal network is composed of T-nodes, each of which represents a time interval during which a clinical event occurred. The starting and ending time instants identify the time interval. Clinical events are organized within a hierarchical structure that corresponds to a model of significant clinical contexts. In an extension of the temporal-network model, each T-node is associated with rules that lead to new conclusions regarding the clinical data.[31]

Recently, a general framework has been proposed for the abstraction of time-stamped and, particularly, clinical data. It is known as the Knowledge-Based Temporal Abstraction (KBTA) method.[11,33] The KBTA framework includes a theoretical model for time and propositions that hold over time; a general inference method; and five specific computational temporal abstraction mechanisms that solve five subtasks, into which the method decomposes the temporal-abstraction task. The output of these mechanisms includes abstractions of *type*, *state*, *gradient*, *rate*, and *pattern*. The five mechanisms require four well-defined types of domain-specific knowledge, the nature of which does not depend respectively on any specific clinical domain but can be specialized for any particular clinical area.[33] The KBTA method had been implemented by the RÉSUMÉ temporal-abstraction system[10] and has been evaluated within several clinical domains, such as oncology, to treat patients who have AIDS, to monitor children's growth, and to manage insulin-dependent diabetes.[11] As temporal primitives, the RÉSUMÉ system uses time stamps at various predefined levels of granularity, typically offset from a clinically relevant time stamp, such as the time of bone-marrow transplantation, the beginning date of chemothera-

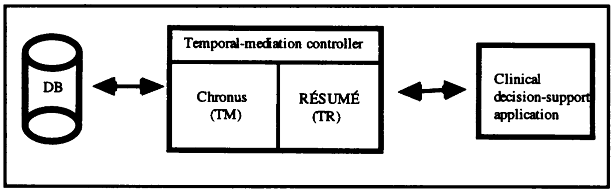

Figure 4.—The Tzolkin temporal-mediation architecture. The Tzolkin mediator enables care providers and decision-support systems to query patient records for complex temporal patterns, possibly involving high-level clinical abstractions. DB=patient electronic database; TM=temporal data-maintenance module; TR=temporal-reasoning module.

py, or the patient's date of birth of the patient (to monitor children's growth). Input data or output abstractions can occur, however, only during time intervals, which are defined as ordered pairs of time stamps. Highly complex patterns can be described and computed, but the set of granularity levels (and therefore the implied temporal uncertainty) is limited to a predefined one that includes minutes, hours, etc.

Absolute-time granularity

The necessity to occasionally provide absolute-time granularity—the capability to refer to the time-axis in multiple ways, not only through different time units—has been addressed by several recent works in medical informatics.[4,6,29] Two different issues must be addressed when providing absolute-time granularity. The first is the representation of uncertainty regarding the location on the time axis of relevant time points or time intervals.[4] The second is the use of time units or references that include not only those associated with the Gregorian calendar, but also ones that are domain-specific (such as the terms *weeks from conception*, *fetal period*, and *infancy*).[4,6,16,35]

One extension to the relational database model that is useful for clinical decision-support applications is the *interval of uncertainty* (Figure 2).[6] Thus, representing a relational database entity that is valid during an interval with indeterminate start and stop instants involves representing explicitly the uncertain start interval, the certain body interval, and the uncertain end interval. Other researchers[4, 37] have described a data model using two different formalisms, both based on an object-oriented approach, that is able to represent intervals and time points given at different and mixed absolute time–granularity, such as the interval referred to in the sentence, "An atrial fibrillation episode occurred on December 14, 1995, and lasted for three minutes."

Calendar-date granularity

Medical applications require systems that are able to represent and manage different time units.[6] Calendar-date temporal granularity is a common one and has been widely studied in the temporal database community.[44]

Conclusions and Future Directions

The temporal-abstraction task and the management of temporal granularity seem to be meeting points between research efforts concerning temporal-reasoning systems and temporal-maintenance systems. An appropriate time model is needed, however, to accomplish each task. Thus, several research issues, most of which are also relevant to general computer science, will be important for the next generation of time-oriented systems in medicine. We mention here only the research avenues that are likely to be most relevant to clinical practice.

1. *Merging the functions of temporal reasoning and temporal maintenance*. By combining these two functions within one architecture, sometimes called a *temporal mediator*, a transparent interface can be added to the patient's database. An example of ongoing research is the Tzolkin temporal-mediation module,[45] which is being developed within the EON guideline-based therapy system.[17] The Tzolkin module combines the RÉSUMÉ temporal-abstraction system,[10] the Chronus temporal-maintenance system,[6] and a controller into a unified temporal-mediation server (Figure 4). The Tzolkin server answers complex temporal queries of health care professionals or clinical decision-support applications, hiding the internal division of computational tasks from the user (or from the clinical decision-support application, which also does not need to have the precise details of the Tzolkin architecture).

2. *Maintenance of both clinical raw data and its abstractions*. Several recent systems allow not only the modeling of complex clinical concepts, but also the maintenance of certain inference operations at the database level. For example, active databases can store and query data that are obtained by the execution of rules triggered by external events, such as the insertion of patient-related data.[46] Furthermore, integrity constraints based on temporal reasoning[25] could be evaluated at the database level, for example, to validate data during their acquisition.

3. *Management of different temporal dimensions of clinical data*. Typically, only the *valid time*, in which the clinical data or conclusions were true, has been considered in medical-informatics research. Also storing the *transaction time*, in which the data were inserted into the patient's record, has multiple benefits, such as being able to restore the state of the database that was true (or what was known) when the physician or a decision-support system decided on a particular therapeutic action. The ability to do so has significance both for explanation and legal purposes. Another temporal dimension of information recently considered is the *decision time*.[47] The decision time of a therapy, for example, could be different from both the valid time, during which the therapy is administered, and the transaction time, at which the data related to the therapy are inserted into the database.

4. *Provision of standardized, user-friendly temporal-query and temporal-visualization interfaces*. Physicians and other health care professionals are not database experts and should not be expected to be familiar with the internal workings of either a temporal-reasoning or a tem-

poral-maintenance system—or with their theoretical underpinnings. Thus, there is a challenge to provide them with easy-to-use, perhaps even graphic, temporal-query interfaces that enable them to take advantage of the sophisticated architectures that are being built on top of the clinical, time-oriented electronic patient records.[48] Many queries might be unnecessary if useful visualization interfaces exist. The semantics of these interfaces, such as deciding automatically which abstraction level of the same set of parameters to show and at what temporal granularity, might draw on the domain-specific knowledge base. An early example was a framework for visualization of time-oriented clinical data,[23] which defined a small but powerful set of domain-independent graphic operators with well-defined semantics, and a domain-specific representation of reasonable temporal granularities for a presentation of various entities in the specific clinical domain. More sophisticated interfaces might be built by taking advantage of, for instance, formally represented knowledge about time-oriented properties of clinical data in specific clinical areas.[11] Indeed, this approach is being taken by the developers of the KNAVE architecture.[49] Using such frameworks, we can build powerful graphical interfaces for visualization of and navigation through multiple levels of abstraction of time-oriented clinical data.

REFERENCES

1. Wiederhold G. Databases for Health Care (LNMI). Heidelberg, Germany: Springer-Verlag, 1981

2. Long WJ, Russ TA. A control structure for time dependent reasoning. *In* Bundy A (Ed): Proceedings of the eighth international joint conference on artificial intelligence. Los Altos, Calif: William Kaufmann, 1983, pp 230–232

3. Pinciroli F, Combi C, Pozzi G, Rossi R. MS2/Cardio—Towards a multi-service medical software for cardiology. Methods Inf Med 1992; 31:18–27

4. Combi C, Pinciroli F, Pozzi G. Managing different time granularities of clinical information by an interval-based temporal data model. Methods Inf Med 1995; 34:458–474

5. Dorda W. Data-screening and retrieval of medical data by the system WAREL. Methods Inf Med 1990; 29:3–11

6. Das AK, Musen MA. A temporal query system for protocol-directed decision support. Methods Inf Med 1994; 33:358–370

7. Dagum P, Galper A, Horvitz E, Seiver A. Uncertain reasoning and forecasting. Int J Forecasting 1993; 11:73–87

8. Dagum P, Galper A. Time-series prediction using belief network models. Int J Human-Computer Studies 1995; 42:617–632

9. Bellazzi R. Drug delivery optimization through bayesian networks—an application to erythropoietin therapy in uremic anemia. Comput Biomed Res 1993; 26:274–293

10. Shahar Y, Musen MA. RÉSUMÉ—a temporal-abstraction system for patient monitoring. Comput Biomed Res 1993; 26:255–273

11. Shahar Y, Musen MA. Knowledge-based temporal abstraction in clinical domains. Artif Intell Med 1996; 8:267–298

12. Larizza C, Berzuini C, Stefanelli M. A general framework for building patient monitoring systems. *In* Barahona P, Stefanelli M, Wyatt J (Eds): Artificial Intelligence in Medicine (Proceedings of AIME '95; LNAI 934). Berlin, Heidelberg: Springer-Verlag, 1995, pp 91–102

13. Rucker DW, Maron DJ, Shortliffe EH. Temporal representation of clinical algorithms using expert-system and database tools. Comput Biomed Res 1990; 23:222–239

14. Keravnou ET, Washbrook J. A temporal reasoning framework used in the diagnosis of skeletal dysplasias. Artif Intell Med 1990; 2:239–265

15. Fagan LM, Kunz JC, Feigenbaum EA, Osborn JJ. Extensions to the rule-based formalism for a monitoring task. *In* Buchanan BG, Shortliffe EH (Eds): Rule-based expert systems—the MYCIN experiments of the Stanford heuristic programming project. Reading, Mass: Addison Wesley, 1984, pp 397–423

16. Keravnou ET. Temporal diagnostic reasoning based on time-objects. Artif Intell Med 1996; 8:235–266

17. Musen MA, Tu SW, Das AK, Shahar Y. EON—a component-based approach to automation of protocol-directed therapy. JAMA 1996; 3: 367–388

18. Jung D, Brennecke R, Kottmeyer M, Hering R, Clas W, Erbel R, Meyer J. Tree-base: a query language for a time-oriented cardiology data base system. *In* Ripley KL (Ed): Computers in Cardiology. Washington, DC: IEEE Computer Society, 1986, pp 381–384

19. Yeh A: Automatically analyzing a steadily beating ventricle's iterative behavior over time. Artif Intell Med 1991; 3:313–323

20. Tong DA, Widman LE. Model-based interpretation of the ECG—a methodology for temporal and spatial reasoning, *In* Safran C (Ed): Annual symposium on computer applications in medical care. New York: McGraw-Hill, 1993, pp 133–139

21. Haimowitz IJ, Kohane IS: Managing temporal worlds for medical trend diagnosis. Artif Intell Med 1996; 8:299–322

22. Bichindaritz I, Conlon E. Temporal knowledge representation and organization for case-based reasoning. *In* Chittaro L, Goodwin S, Hamilton H, Montanari A (Eds): Third International Workshop on Temporal Representation and Reasoning (TIME '96). Los Alamitos, Calif: IEEE Computer Society Press, 1996, pp 152–159

23. Cousins SB, Kahn MG. The visual display of temporal information. Artif Intell Med 1991; 3:341–357

24. Chitarro L, Dojat M. Using a general theory of time and change in patient monitoring—experiment and evaluation. Comput Biol Med 1997; 27:435–452

25. Horn W, Miksch S, Egghart G, Popow C, Paky F. Effective data validation of high-frequency data—time-point, time-interval, and trend-based methods. Comput Biol Med 1997; 27:389–409

26. Larizza C, Moglia A, Stefanelli M. M-HTP—A system for monitoring heart transplant patients. Artif Intell Med 1992; 4:111–126

27. Degoulet P, Devries C, Chantalou J, Klinger E, Sauquet D, Zweigenbaum P, Aimé F. LIED—A temporal database management system. *In* Salamor R, Blu B, Mogens J (Eds): MEDINFO '96—Proceedings of the Fifth Conference in Medical Informatics. Amsterdam: North-Holland, 1986, pp 532–536

28. Aliferis CF, Cooper GF. A new formalism for temporal modeling in medical decision-support systems. *In* Gardner RM (Ed): Annual Symposium on Computer Applications in Medical Care. Philadelphia, Penn: Hanley & Belfus, 1995, pp 213–217

29. Dolin RH. Modeling the temporal complexities of symptoms. J Am Med Inform Assoc 1995; 2:323–331

30. Kohane IS. Temporal reasoning in medical expert systems. Technical Report 389, Laboratory of Computer Science, Massachusetts Institute of Technology: Cambridge, Mass, 1987

31. Kahn M, Fagan L, Tu S: Extensions to the time-oriented database model to support temporal reasoning in medical expert systems. Methods Inf Med 1991; 30:4–14

32. Allen J. Towards a general theory of action and time. Artif Intell 1984; 23:123–154

33. Shahar Y. A framework for knowledge-based temporal abstraction. Artif Intell 1997; 90:79–133

34. Goralwalla IA, Ozsu MT, and Szafron D. Modeling medical trials in pharmacoeconomics using a temporal object model. Comput Biol Med 1997; 27 (5):369–387

35. Shahar, Y, Miksch, S, Johnson, PD. An intention-based language for representing clinical guidelines. *In* Proceedings of the 1996 AMIA Annual Fall Symposium (formerly the Symposium on Computer Applications in Medical Care), Philadelphia, Penn: Hanley & Belphus, 1996:592–596

36. Combi C, Pinciroli F, Pozzi G. Temporal clinical data modeling and implementation for PTCA patients in an OODBMS environment. *In* Proceedings of Computers in Cardiology. Los Alamitos, Calif: IEEE Computer Society Press, 1994:505–508

37. Combi C, Cucchi G, Pinciroli F. Applying object-oriented technologies in modeling and querying temporally oriented clinical databases dealing with temporal granularity and indeterminacy. IEEE Trans on Inform Tech in Biomed 1997; 1:100–127

38. Shortliffe EH, Perreault LE (Eds); Wiederhold G, Fagan LM (Asst Eds). Medical informatics—Computer applications in health care. Reading, Mass: Addison Wesley, 1990

39. Blum RL: Discovery and representation of causal relationships from a large time-oriented clinical database (LNMI). New York: Springer Verlag, 1982

40. Console L, Torasso P. Temporal constraint satisfaction on causal models. Inf Sciences 1993; 68:1–32

41. Kouramajian V, Fowler J. Modeling past, current, and future time in medical databases. *In* Ozbolt JG (Ed): Annual symposium on computer applications in medical care. Philadelphia, Penn: Hanley & Belfus, 1994, pp. 315–319

42. Snodgrass R, Ahn I. Temporal databases. IEEE Computer 1986; 19:35–42

43. Kahn M, Fagan L, Tu S. Combining physiologic models and symbolic methods to interpret time-varying patient data. Methods Inf Med 1991; 30:167–178

44. Wiederhold G, Jajodia S, Litwin W. Dealing with granularity of time in temporal databases. *In* Proceedings of the third Nordic conference on advanced information systems engineering. Trondheim, Norway, 1991, pp 124–140

45. Nguyen J, Shahar Y, Tu SW, Das AK, Musen MA. A temporal database mediator for protocol-based decision support. *In* Proceedings of the 1997 AMIA. Annual Fall Symposium (formerly the Symposium on Computer Applications in Medical Care). Philadelphia, Penn: Hanley & Belfus, 1997, pp 298–302

46. Caironi PV, Portoni L, Combi C, Pinciroli F, Ceri S. HyperCare: A prototype of an active database for compliance with essential hypertension therapy guidelines. *In* Proceedings of the 1997 AMIA Annual Fall Symposium (formerly the Symposium on Computer Applications in Medical Care). Philadelphia, Penn: Hanley & Belfus, 1997, pp 288–292

47. Gal A, Etzion O, Segev A. Representation of highly complex knowledge in a database. J Intell Info Sys 1994; 3:185–203

48. Combi C, Pinciroli F, Musazzi G, Ponti C. Managing and displaying different time granularities of clinical information. *In* Ozbolt JG (Ed): The 18th Annual Symposium on Computer Applications in Medical Care. Philadelphia, Penn: Hanley & Belphus, 1994, pp 954–958

49. Shahar Y, Cheng C. Ontology-driven visualization of temporal abstractions. *In* Proceedings of the Eleventh Workshop on Knowledge Acquisition, Modeling, and Management. Alberta, Canada: Banff, 1998 (in press)